Who am I? Where am I?

Geoff Huston
May 2001

Well, it may have been a pretty tough weekend for me, but I assure you that I can still remember everything. The question posed was not personal, but was intended to start us on an examination of the role of an Internet address and its implications for IP mobility.

An IP address has a heavy burden to carry. Not only does it uniquely identify an Internet-connected device (*who*), but it also used to describe the location of the device within the network topology (*where*), and it's also used as the basis of determining a path to that location (*how*). This combination of who, where and how, embedded into each and every IP address, is quite a load for a single protocol value. While this combination of roles has served us well for some decades of Internet growth, it is appropriate to review this approach as we move on. We are now transitioning from a network with a dimension measured in units of millions to one with dimensions measured in units of billions: billions of connected devices, serving billions of packets per second. As we scale the Internet it is appropriate to review the current design parameters and ensure that will still carry us forward into larger and larger scales of networks. One of the more fundamental design decisions in IP is the joining of identity, location and path into a single IP address. Are we getting close to that point when we will have to unbundled this combination, and treat identity, location and path as distinct concepts within the IP architecture?

The use of a distinct identity and location is a very old one in communications systems, and dates back centuries, if not millennia. The postal system copes quite adequately, most of the time, using envelope labels, which contain an identity field and a hierarchically structured delivery address, or location. If you change your postal address you get to keep your identity, and you would accept nothing less. But when you unplug your Internet-connected laptop from its socket on your desk, take it elsewhere and then use a dial-up connection to access the Internet, your laptop has acquired a new IP address. With this new IP address is an association of an entirely new identity. Each time you move, or each time you dial, you get a new IP address and with that new address is a new identity.

Now for an Internet which is dominated by client-server web applications, and where the servers are fixed, and the identity of the client is not critical to the server, this may well represent an acceptable compromise. But if we've learned one thing from the past decade of Internet growth, it's that when you construct an entire network around the characteristics of a particular class of application, then a change of application ultimately requires an entirely new network. If you want a network infrastructure to have some longevity it must avoid making application-based assumptions within its design. So while client-server may dominate today's Internet application environment, its unwise to believe that this is the only application that the Internet should support. Equally it's unwise to assume that fixed devices will dominate the Internet. Projections by a number of mobile device vendors put the number of mobile IP devices as greater than the number of conventional fixed devices by sometime in 2003. While the date may be debatable, this projection heralds a very fundamental change for the architecture of the Internet. What we appear to be looking at over the next couple of years is a dramatic increase in the number of Internet connected devices, coupled with an emphasis on service to mobile and roaming devices, and a potential shift away from a server-client application architecture to one which makes more use of identity-based services as well as peer-to-peer. In all of these changes, the concepts of identity and location are critical.

One of the fundamental tasks of scaling in a mobile and roaming environment is to use the approach of a fixed identity while allowing the location to vary. If the end-to-end session level

transactions use the identity fields as the reference points for the session, then the underlying mobility of the end devices will be transparent to the application session. For the packets to pass through the network, the device's identity must be associated with a current location. As the end device roams through the network, the association of identity to location needs to be updated.

InIPv4 the approach to mobility has been to use two IP addresses to fully describe a mobile device. The mobile device uses a constant IP address, which in this case can be considered as its identity. This identity IP address is passed to the local mobile base station, who then informs the mobile's home station of the mobile's identity address, as well as the address of the current mobile base. Any packets that the home station want to pass to the mobile device can be sent to the mobile base station, who in turn will pass them on. From this perspective, the address of the mobile base station can be thought of as the mobile device's current location, while its own IP address serves as an identity. The mode of operation of mobility in this model is quite interesting. A remote system sends a packet to the mobile device quite normally – that is using the mobile device's IP address as both an identity and location identifier. Once the packet reaches the home base station the packet is encapsulated in an IP transport header, with the new destination IP address being that of the most recent mobile base station. In effect, the packet now uses a location IP address that is not the same as its identity IP address. At the mobile base station the outer IP transport header is stripped off and the original packet is passed directly to the mobile device.

This approach of adding an outer IP transport header to delineate location from identity can be considered somewhat cumbersome. One approach being considered within the continued evolution of IPv6 is to divide the 128-bit address into two parts, a 64 bit routing segment and a 64-bit identity segment. This approach has the potential to allow a mobile device to maintain a constant identity in the low order 64 bits of its IPv6 address, and use a location-defined high order 64-bit value as its current location. While there is still a fair amount of further refinement that must take place with this approach, it appears to be a solid step towards recognizing that as we transform the Internet into a larger and predominately mobile Internet we will need to support mobility and roaming. Allowing a device to maintain a constant identity in one part of an IP address while allowing its location to determine the value of the other half of the address appears to offer a way to support seamless mobility to the extent that the application level need not be aware that the device is roaming or fixed.

Of course having a persistent identity offers more than support for mobility in IP. Currently, the IPSEC end-to-end security protocol uses the host's IP address as an integral part of the security association. Having a persistent and unique identity allows a security association to be based on the identity of the two hosts, rather than on the location-based IP addresses they happen to be using at the time. Persistent identity can also be used to support multi-homing. If two hosts in multi-homed networks establish a session using the identity parts of the address fields as the session identifiers, then it is feasible that they can swap location parts of their addresses in response to network failure and still maintain session integrity.

A word of caution is abut privacy is necessary. It is often desirable to be able to initiate a communication without revealing your identity. Placing an identity value in the source header of every packet your system sends compromises your ability to keep your identity private when you wish to initiate an anonymous communication. Paradoxically, it appears that together with the need to have a persistent and unique identity embedded at the IP level, there is also a need to allow this identity field to be masked out on demand.

Much of the leverage of the Internet lies in taking functionality that has traditionally been viewed as part of a network's role and passing that function to the end devices. On the whole this makes for simpler, cheaper and more efficient networks, and makes for hosts, which are capable of adapting to the current operating state of the network. Allowing an Internet host to have a clear and persistent view of its own identity and allowing a host's location to reflect its

current position within the topology of the network can be see as yet another step in providing hosts with greater levels of capability and adaptability, particularly in an Internet which is poised to become the mainstay of mobile data services.